

# Mutual information for fitting deep nonlinear models

Jacob S. Hunter<sup>1</sup> and Nathan O. Hodas <sup>\*1</sup>

<sup>1</sup> *Pacific Northwest National Laboratory*  
*{jacob.hunter, nathan.hodas}@pnnl.gov*

December 20, 2016

## Abstract

Deep nonlinear models pose a challenge for fitting parameters due to lack of knowledge of the hidden layer and the potentially non-affine relation of the initial and observed layers. In the present work we investigate the use of information theoretic measures such as mutual information and Kullback-Leibler (KL) divergence as objective functions for fitting such models without knowledge of the hidden layer. We investigate one model as a proof of concept and one application of cognitive performance. We further investigate the use of optimizers with these methods. Mutual information is largely successful as an objective, depending on the parameters. KL divergence is found to be similarly successful, given some knowledge of the statistics of the hidden layer.

## 1 Introduction

There is extensive literature on the effects of sustained activity on human cognitive performance, which suggests that performance is determined by some finite resource that causes a decline in performance when depleted [1, 2, 10]. However, the nature of this resource remains poorly understood. In other work (to be submitted) we propose a model in which standardized test performance depends on cognitive resources that are depleted as questions are answered. In this work, we present a method for fitting ODE's to data when the mapping between the observables and underlying hidden variables is unknown. For the cognitive performance example, a researcher may propose a microscopic mechanistic model for resource variation in the brain but not know how the variations in the physiological resources translate into variation in the observed performance.

---

\*Corresponding author

We can characterize the relation of cognitive performance to physiological resource usage as a deep nonlinear system. We consider a nonlinear system to be deep if it can be represented as a composition of multiple nonlinear transformations. There is an initial hidden layer, followed by a transformation layer, and an output layer that corresponds to what is actually observed. In short, complex nonlinear dynamics on the microscopic scale ultimately translate to emergent behavior on the macroscale via multiple scales, spatially and temporally. Many types of objective functions may be used to fit equations to hidden variables, such as Kalman filters, Bayesian models, and simple root-mean-squared error (RMSE) regression via supervised learning[8, 9, 12]. In general, these models are parametric, and their success requires some fundamental assumptions about the dynamics producing data. For example, to minimize RMSE, one needs to be able to regress to the final observed variable, so the hidden layer needs to also be known or learned.

One approach to understanding a deep nonlinear system would be to propose a model for the hidden variables, construct a function approximator (such as a neural network) as the transformation layer, and then use a fitting procedure such as backpropagation[8] to jointly regress the generic transformation layer and fit the underlying hidden model. The viability of this approach depends on the fidelity of the function approximator, which may require an extremely large amount of data to properly determine, akin to traditional deep learning. Here, we propose a method that is agnostic to the transformation layer by seeking to maximize explanatory power of the underlying hidden model to the ultimate observables.

We investigate information theoretic measures, such as mutual information and Kullback-Leibler divergence, as objective functions for fitting the model to generated data. We construct a number of example nonlinear systems to validate the ability for these information theoretic measures to ‘peer through’ the intervening nonlinear unknown layers and directly fit observations to the hidden variables we are attempting to model. We confirm that mutual information may be a valuable tool for fitting models to deep nonlinear systems.

## 1.1 Deep Nonlinear Models

In the present work, we consider systems that we may call “deep nonlinear models,” in that we may have a clear hypothesis for what the dynamics of the initial layer may be, but there remains an unknown layer that transforms the output of the dynamics into what is observed. We may consider deep nonlinear models to take the form

$$\dot{x} = f(x, t) \quad \text{Initial Layer} \quad (1)$$

$$\dot{y} = g(x, y, t) \quad \text{Hidden Layer} \quad (2)$$

$$z = h(y, t). \quad \text{Observable/Measurement Layer} \quad (3)$$

where  $f, g, h$  are arbitrary transformations.

This hidden layer may be due to some unknown dynamics, stochasticity, etc. For example, consider the task of using the number of people at the beach to infer the weather. We may know how to predict the weather (initial layer), and we may separately have an empirical count of beach-goers (observation), but the transformation from the weather to how people decide when to visit the beach remains hidden. Because we only care about making estimates of the weather, we would prefer to not have to model people’s decision making processes explicitly.

Frequently, we may only have the observations – and not even know the initial layer. For example, we may wonder if we can count people at the beach to estimate the weather at different locations. More generally, we may attempt some sort of model identification given some limited sensor activity. To be successful, we need to ‘see through’ the hidden layers to directly fit the parameters of the initial layer. Many techniques exist to do this. Here, we present an information theoretic technique based on mutual information that allows us to non-parametrically conduct fitting under weak assumptions about the dynamics of the hidden layer.

## 1.2 Mutual Information

Mutual information ( $MI$ ) quantifies how many bits of entropy we may reduce our uncertainty in  $X$  by knowing another variable  $Y$  and vice-versa. For example, how much information about an individual’s performance ( $X$ ) do we obtain by knowing the amount of cognitive resources available to him or her ( $Y$ )?  $MI$  is defined as  $MI(X, Y) = H(X) + H(Y) - H(XY)$ , where  $H(X)$  is the entropy of the random variable  $X$ , and  $H(XY)$  is the entropy of the joint distribution of  $X$  and  $Y$ . In the present work, we leverage the property that  $MI(X, Y) = MI(f(X), g(Y))$  for invertible functions  $f$  and  $g$  [7], so the underlying initial layer of the nonlinear system may be fit from the observations as long as the two can be related via an isomorphism. If the relations between  $X$  and  $Y$  are not isomorphic, we may still utilize mutual information, but we will only be able to explain limited portions of the dynamics. In contrast, optimizing a regression model or Pearson correlation for parameter estimation requires not only correctly modeling the initial layer but also knowing how those resources quantitatively translate into the observations, because  $R^2$  is only maximized when predicted and observed results have an affine relation.

For example, in Fig. 1, we demonstrate how the mutual information between two variables remains unchanged under isomorphisms. However, the correlation between the transformed variables may go from positive to zero to negative, depending on the transformation. Only the cosine transformation, which is not an isomorphism, shows a decrease in mutual information.

This ability to quantitatively relate the explanatory power of two variables, regardless of the isomorphic transformation between them, allows this method to relate the dynamics of the initial layer with the observation in the final layer. In addition, mutual information

is well defined for discrete, continuous, or categorical data – or some combination thereof. In this work, we present a method for using mutual information to fit dynamics of the initial layer to discrete observations, without having to know the intervening hidden layer.




Figure 1: Mutual information does not vary under invertible transformations, while correlation is only maximize with a perfect linear relationship.

### 1.3 KL Divergence

In addition to mutual information, we tested the Kullback-Leibler (KL) divergence as an objective function. KL divergence is defined as the amount of information lost when one distribution is approximated by a second; it can thus be thought of as a distance between the two distributions, although it is not symmetric. The KL divergence of continuous random variables from  $Q$  to  $P$  is given by

$$D_{KL}(P||Q) = \int_{-\infty}^{\infty} p(x) \ln \frac{p(x)}{q(x)} dx \quad (4)$$



When testing KL divergence as an objective function, we considered the distribution of observations against our initial model under a variety of comparisons. First, we considered the divergence from the fitted model to the resource values used to generate the outcome data, given the outcome. In this case, the objective is a loss function, as the goal is for the model to closely approximate reality. Finally, we considered the divergence from the distribution of resource values given one outcome to the distribution given the other outcome; when this divergence is maximized, it forces the distributions to be as different as possible, which should increase the explanatory power of the model.

## 2 Approach

### 2.1 Example with Mutual Information

As an example and proof of concept of our fitting with mutual information, we construct a family of deep nonlinear models:

$$x = e^{-\lambda t} \tag{5}$$

$$y = x + y_{\text{err}} \tag{6}$$

$$z = f(y), \tag{7}$$

with  $y_{\text{err}}$  being Gaussian noise and  $z = f(y)$  having one of three functional forms:

$$f_1(y) = ay \tag{8}$$

$$f_2(y) = e^{ay} \tag{9}$$

$$f_3(y) = \sin(ay). \tag{10}$$

In each case, we wish to fit  $\lambda$  to the generated data given the time series of  $x$  and  $z$ . To generate the data, we chose a value of  $\lambda$  and a range for  $t$ . With a fixed  $\lambda$  and a given  $t$ , we can then easily determine  $f_1$ ,  $f_2$ , and  $f_3$ .

To evaluate if a maximum in mutual information occurs with the true  $\lambda$ , we simply plotted the  $MI$  between data generated from  $\lambda$  and different values of  $x$  generated from a series of values of guesses  $\hat{\lambda}$ : we varied  $\hat{\lambda}$  in increments of 0.01 in the range (-1, 4) and calculated the mutual information  $MI(z, e^{-\hat{\lambda}x})$ . This entire process was repeated for several values of  $\lambda$  and  $a$ . Because of the highly deterministic relationship between  $x$ ,  $y$  and  $z$ , we estimated mutual information using the Non-parametric Entropy Estimation Toolbox (NPEET) with Local Non-uniformity Correction (LNC)\* [4, 6]. This produced far more robust, less biased estimates of mutual information, by correcting for systematic errors that can emerge for deterministic data for the Kraskov estimator.

Fitting the example model using mutual information correctly produced a maximum of the objective function at or very close to the true value for  $\lambda$  (Fig. 2). In particular,

---

\*[https://github.com/BiuBiuBiLL/NPEET\\_LNC](https://github.com/BiuBiuBiLL/NPEET_LNC)

Fig. 2(c) shows a sharp maximum at the true value of  $\lambda = 2$  for all forms of  $z$ . While the mutual information curves for the exponential and sinusoidal  $z$  in Figs. 2(a) and 2(b) show smoother maxima that are not perfectly at the true value, they provide a fairly accurate estimate of the true value. The dip in mutual information to 0 at  $\lambda = 0$  is due to  $y$  taking the form  $e^{0x} = 1$ ; in practice, this would also prevent a local optimizer from transitioning over regimes when a particular parameter completely decouples the hidden model from the observations. It would require a more global optimization strategy.

## 2.2 Application: Cognitive Performance

A particular system we consider in the present work is cognitive performance over time. Previous work established a model for cognitive performance based on lactate metabolism [3, 15, 16]. In this model, performance is determined by a primary resource,  $A$ , which is low during normal resting conditions. Work on a task consumes  $A$ . However,  $A$  is replenished by conversion of a secondary resource,  $B$ , into  $A$ ;  $B$  then recovers during resting conditions. These dynamics are modeled as:

$$\dot{A}(t) = w(A, B, t) - \frac{k_w}{t^\rho} \frac{A(t)}{K_A + A(t)} \quad (11a)$$

$$\dot{B}(t) = -w(A, B, t) + \frac{k_r}{t^\rho} (B_{max} - B(t)) (1 - \delta(t)) \quad (11b)$$

$$w(A, B, t) = \frac{k_b}{t^\rho} \frac{(1 - A(t)) B(t)}{K_B + B(t)} \delta(t), \quad (11c)$$

where  $w$  is the rate of conversion of  $B$  into  $A$  and  $\delta(t)$  is given by 1 when on task (working) and 0 when off task (resting).


Although this model may be difficult to understand outside the context of modeling cognitive performance, it captures many of the properties that make fitting difficult for many methods. Clearly, this model is complex, nonlinear and time-heterogeneous. In addition, it is non-differentiable at the boundaries between tasks. Furthermore, although this may be a proposed model of cognitive resources, we have no hypothesis for how cognitive resources translate into correct/incorrect answers on a standardized test. All we can assume is that fewer resources decreases the likelihood of correct answers. Thus, to fit the model, we need to be ambivalent to the intervening hidden layer.

## 2.3 Generated Performance Data


To test fitting the model, we generated data similar to data from online practice standardized tests, obtained from the Kaggle “What do you know” competition<sup>†</sup> and provided

---


<sup>†</sup><http://www.kaggle.com/c/WhatDoYouKnow>



(a)  $a = 1, \lambda = 0.5$



(b)  $a = 2, \lambda = 1$



(c)  $a = 3, \lambda = 2$

Figure 2: Mutual information of outer layer  $z = f(y)$  with fitted hidden layer  $y' = e^{-\lambda x}$  for different forms of  $f(y)$  and true values of  $\lambda$ . True value for  $\lambda$  is given by dashed line.

by grockit.com. Problem times were generated stochastically from an exponential distribution,

$$p(x) = \lambda e^{-\lambda x}, \quad (12)$$

with  $x$  the time spent on or off task, in minutes. For time on task, we set  $\lambda = \frac{1}{4}$ ; for time off task,  $\lambda = \frac{1}{4}$ , except for every tenth task, when  $\lambda = \frac{1}{40}$ . Our fictional participant thus spent an average of four minutes on or off each task, with a longer break every ten tasks. To generate the dataset, 25 time series with 3000 event times were generated in this fashion. From the generated time series, we used a numerical ODE solver from the `scipy.integrate` library to get the primary and secondary resources from the model using the previously obtained parameters. Outcome of a task was determined randomly based on the primary resource,  $A$ , at the end of each question; the probability of success was given by the logistic sigmoid function,

$$P(\text{correct}|A) = (1 + e^{-\alpha(A-A_0)})^{-1}, \quad (13)$$

with  $\alpha$  and  $A_0$  chosen as 169 and 0.204, respectively, such that typical high values of  $A$  would give a success rate of 70% and typical low values would give a success rate of 30%.

Using this generated data, we sought to refit the resource model to determine whether our fitting methods are appropriate. To start, we fit each parameter individually to get a sense for the shape of the objective function. To fit a parameter, we varied around the true value and solved the model ODEs for the resources at each time. We then used the newly calculated model resources and the outcome data to calculate the objective function at each parameter step.

## 2.4 Objective Functions

To construct an objective function using mutual information, we estimated mixed continuous-discrete mutual information using the implementation based on Kraskov et. al. and Kozachenko and Leonenko in NPEET<sup>‡</sup> [6, 5, ?]. We calculated the mutual information of the series of task outcomes (with 1 standing for correct/success and 0 for incorrect/failure) with the primary resource at the end of each task.

KL divergence was also estimated using the implementation in NPEET [14, 5]. We used two methods for using KL divergence as an objective function. In both cases, we considered the frequency distribution of the primary resource at the end of a task given the outcome. The first use of KL divergence relied upon prior knowledge of the underlying resources used to generate the data. We again used the distribution of fitted resources given outcome, but now compared to the distribution of the underlying resources:

$$D_{KL}(p(A_{\text{fitted}}|\text{correct})||p(A_{\text{underlying}}|\text{correct}))$$

---

<sup>‡</sup><http://www.isi.edu/~gregv/npeet.html>

and

$$D_{KL}(p(A_{\text{fitted}}|\text{incorrect})||p(A_{\text{underlying}}|\text{incorrect})).$$

In the second method, we compared the frequency distribution of the primary resource given a successful outcome with that given an unsuccessful outcome:

$$D_{KL}(p(A|\text{correct})||p(A|\text{incorrect}))$$

and

$$D_{KL}(p(A|\text{incorrect})||p(A|\text{correct})).$$

By maximizing either of these divergences or a combination thereof, the model should be as distinct as possible within a regime of cognitive depletion compared to a regime of plentiful cognitive resources.

## 2.5 Optimizers

Because mutual information does not require the random variables to be positively correlated for it to be maximized, using it as a metric can result in model resource values that are lower on average for successful outcomes than for unsuccessful outcomes. While this could be interpreted as the model representing waste products instead of resources, it goes against the assumptions that base the model on lactate metabolism. Thus, we added a constraint that the average resource given success be greater than that given failure:

$$f_c = \langle A|\text{correct} \rangle - \langle A|\text{incorrect} \rangle \geq 0, \quad (14)$$

where  $A$  is the primary resource modeled in Eq. 11.

We ran tests on optimizing all parameters together over the mutual information objective using constrained optimizers implemented in NLOPT<sup>§</sup>. However, the optimizers here posed problems for this task. Because the objective function requires solving the model ODEs for the time series and parameters, it takes a long time to evaluate. Additionally, because the Kraskov mutual information is non-differentiable, the gradient is unknown and must be estimated; therefore, each iteration of the optimizers required multiple evaluations of a computationally expensive objective. Thus, we use Simultaneous Perturbation Stochastic Approximation (SPSA). SPSA lends itself to objective functions requiring a simulation or other computationally expensive task by requiring only two calls of the objective function to estimate the gradient[11, 13].

## 3 Discussion

When fitting one parameter from the true value and holding the rest constant, mutual information as an objective function produced a maximum at or near the correct param-

---

<sup>§</sup>Steven G. Johnson, The NLOpt nonlinear-optimization package, <http://ab-initio.mit.edu/nlopt>

eter values for  $k_w$ ,  $k_r$ , and  $B_{max}$ ; the landscapes for  $\rho$  and  $k_b$  featured a plateau-like range of values at the respective maxima (Fig. 3).

KL divergence comparing distribution of resource values to an *a priori* distribution produced a minimum at the true parameter value for most of the parameters. However,  $\rho$  did not have its minimum near the true value, but at a greater value. Similarly to the mutual information analysis,  $k_b$  had an optimum value for a range of values with the true value near the lower end (Fig. 4).

KL divergence comparing the distributions of resource values given outcome fails to produce a maximum at the true parameter; in fact, it produces a local minimum near the true value for  $k_w$ ,  $k_r$ , and  $B_{max}$ . The parameter  $k_b$  was the only parameter to produce a maximum at the true value, but similarly to its behavior with other objective functions, it also produced optima at larger values (Fig. 5).

While tests with KL divergence comparing the model to the original resource values produced the clearest optimum at the true parameter values, it poses a problem that the distribution of resource values at the end of each task must be known *a priori*. In the case of the generated data, the distribution of resource values was found to be Gaussian, allowing the divergence to be approximated using an approximation for the true distribution; however, even if the functional form of the true distribution is known *a priori*, its parameters may not be. On the other hand, using KL divergence as an objective function to encourage disjointness of resource probability distributions given outcome failed to recover the underlying parameters of the model. It is worth noting that both successful fitting methods have poor sensitivity for fitting  $\rho$ . The unusual behavior in the fitting of  $\rho$  and  $k_b$  seems to indicate that they are less important to the overall fit of the model, except when  $k_b$  is too low. This may indicate that establishing a lower bound may be more important when fitting this parameter.

## 4 Conclusion

In this work we investigated fitting the hidden layer of deep nonlinear models using mutual information and KL divergence. Although this was a preliminary investigation, we demonstrated the potential value of using mutual information to fit deep nonlinear models. Although it may be sensitive to biases inherent in the underlying mutual information estimators, the property that mutual information is invariant under isomorphic transformations allows it to estimate the sensitivity of an observation to a model when there is an unknown hidden layer between them. When using KL divergence to compare a model to prior-known statistics of the hidden variables, true parameter values are similarly recovered. However, using KL divergence to maximize explainability by maximizing the divergence between frequency distributions of the hidden layer given the observed layer completely fails in recovering the original parameter values.

Future work will further explore the formal relationship between mutual information

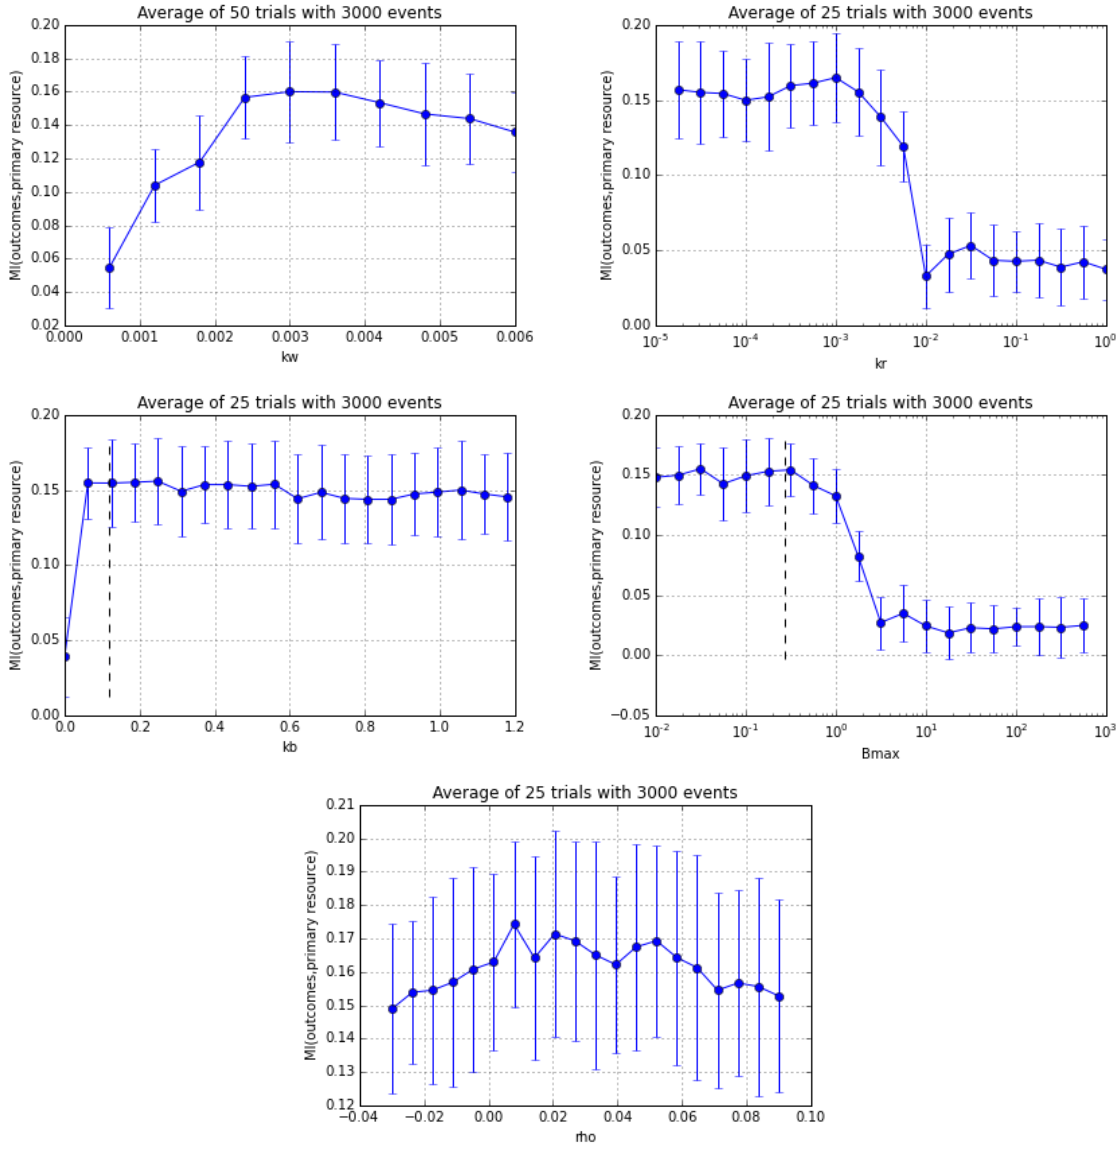


Figure 3: Mutual information of primary resource with task outcomes. True parameters are given by dashed lines.

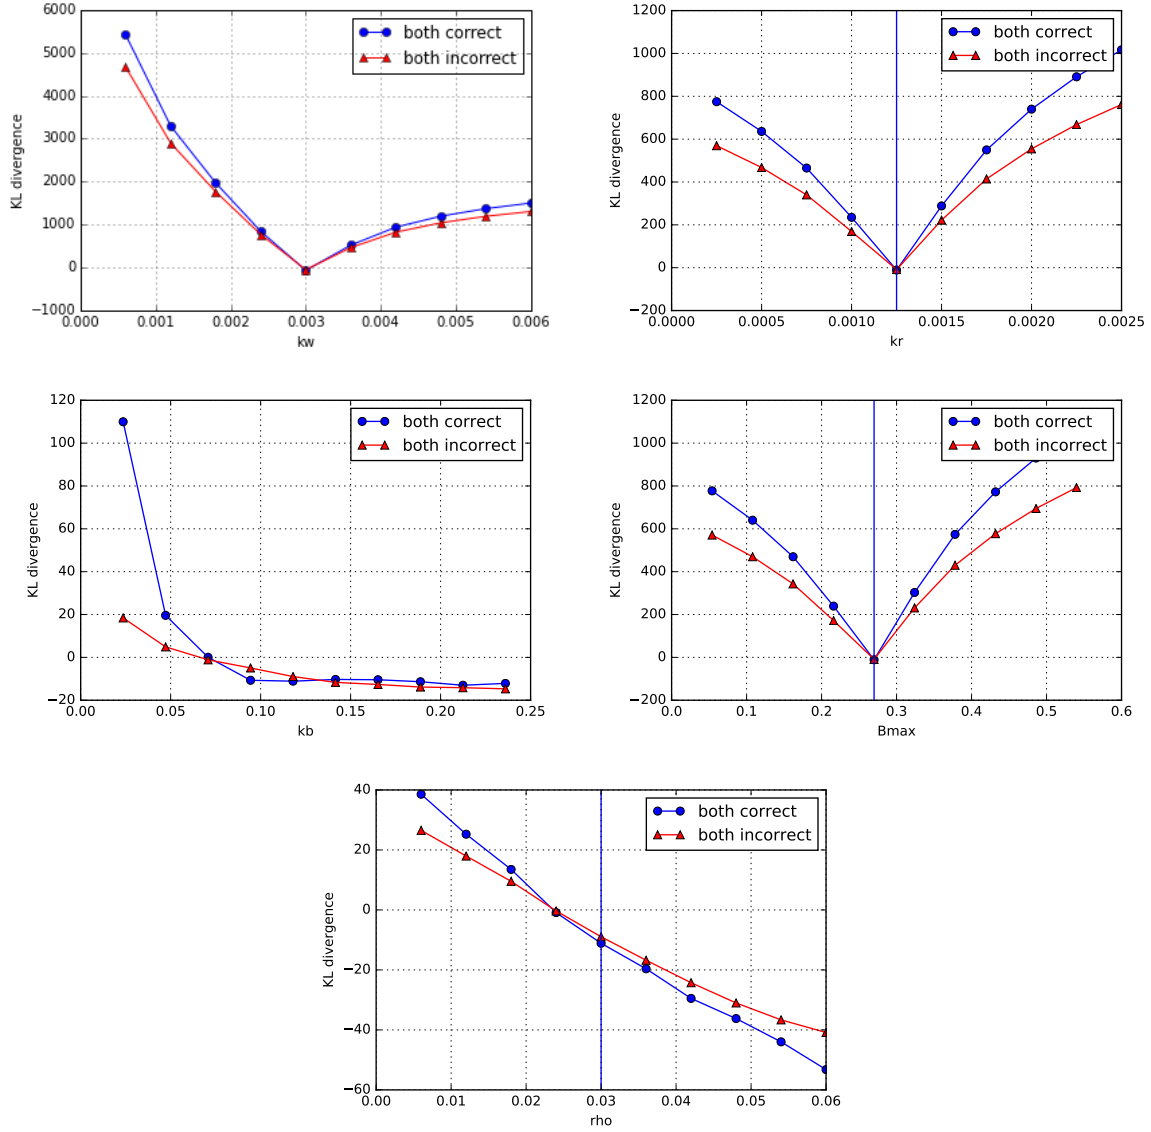


Figure 4: KL divergence of frequency distributions of resource from the fitted model with that of the original hidden resources. True parameters are given by the vertical lines.



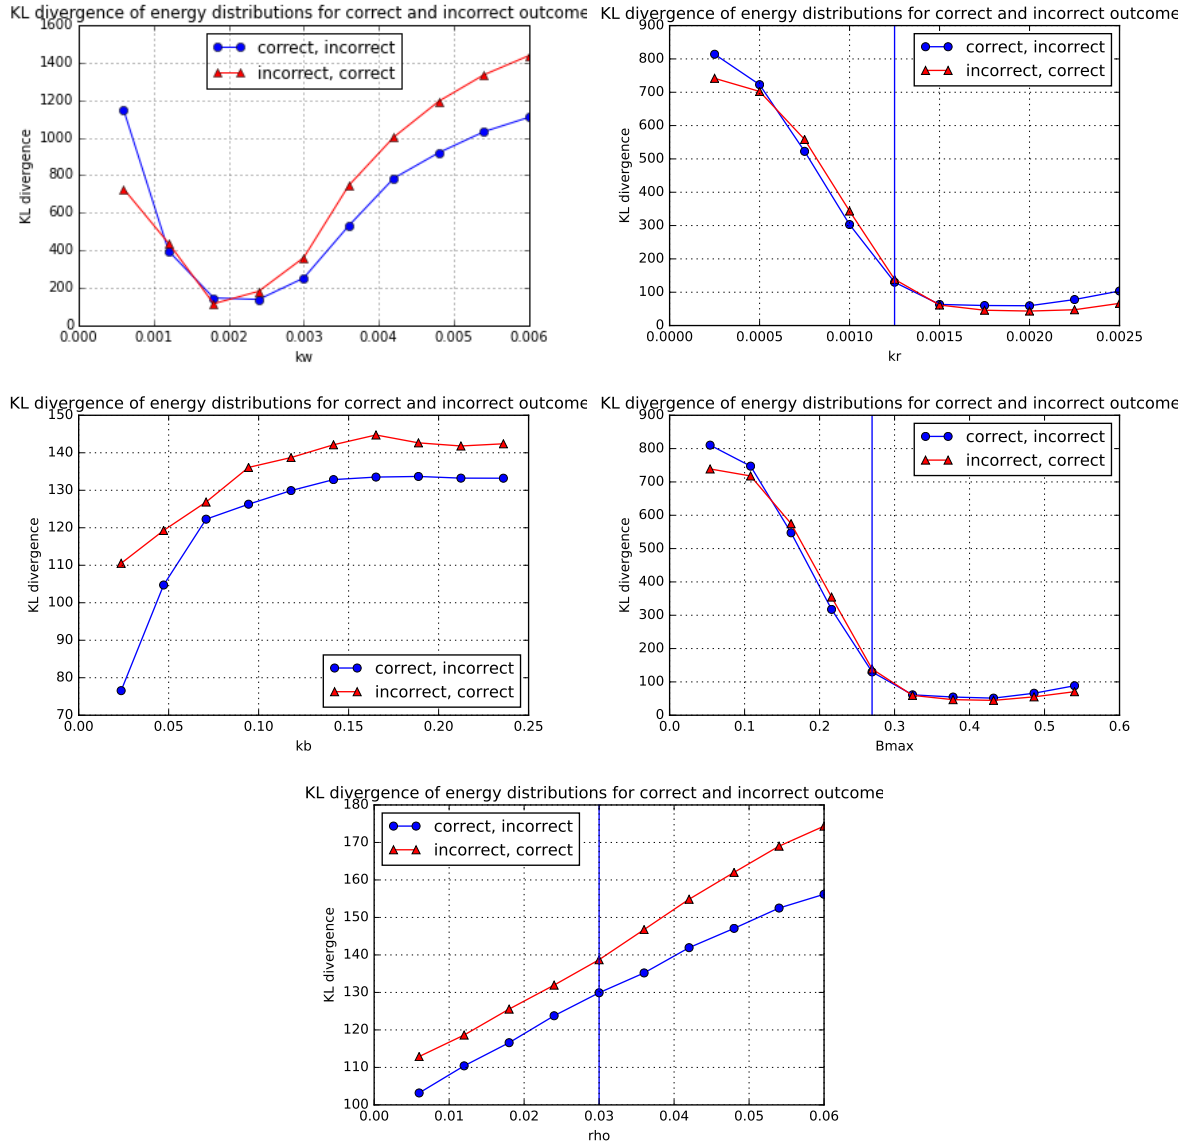


Figure 5: KL divergence of frequency distributions of resource from the fitted model with that of the original hidden resources. True parameters are given by the vertical lines.

and variations in parameters in the initial, hidden, and observation layers of deep nonlinear models. The present work provides a promising proof of principle in simple and complex models. The applications for this work include cognitive science, system identification, and machine learning.

## 5 Acknowledgments

This work was funded by the Analysis in Motion Initiative at Pacific Northwest National Laboratory.

## References

- [1] M. A. Boksem, T. F. Meijman, and M. M. Lorist. Effects of mental fatigue on attention: an erp study. *Brain Res Cogn Brain Res*, 25(1):107–16, 2005.
- [2] M. A. Boksem and M. Tops. Mental fatigue: costs and benefits. *Brain Res Rev*, 59(1):125–39, 2008.
- [3] M. Cloutier, F. B. Bolger, J. P. Lowry, and P. Wellstead. An integrative dynamic model of brain energy metabolism using in vivo neurochemical measurements. *J Comput Neurosci*, 27(3):391–414, 2009.
- [4] Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. Efficient estimation of mutual information for strongly dependent variables. *CoRR*, abs/1411.2003, 2014.
- [5] S. Khan, S. Bandyopadhyay, A. R. Ganguly, S. Saigal, 3rd Erickson, D. J., V. Protopopescu, and G. Ostrouchov. Relative performance of mutual information estimation methods for quantifying the dependence among short and noisy data. *Phys Rev E Stat Nonlin Soft Matter Phys*, 76(2 Pt 2):026209, 2007.
- [6] A. Kraskov, H. Stogbauer, and P. Grassberger. Estimating mutual information. *Phys Rev E Stat Nonlin Soft Matter Phys*, 69(6 Pt 2):066138, 2004.
- [7] Chandra Nair, Balaji Prabhakar, and Devavrat Shah. On entropy for mixtures of discrete and continuous variables. *CoRR*, abs/cs/0607075, 2006.
- [8] Oliver Nelles. *Nonlinear system identification: from classical approaches to neural networks and fuzzy models*. Springer Science & Business Media, 2013.
- [9] V Peterka. Bayesian system identification. *Automatica*, 17(1):41–53, 1981.
- [10] A. K. Shah, S. Mullainathan, and E. Shafir. Some consequences of having too little. *Science*, 338(6107):682–5, 2012.

- [11] J. C. Spall. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *Ieee Transactions on Aerospace and Electronic Systems*, 34(3):817–823, 1998.
- [12] Henning U Voss, Jens Timmer, and Jürgen Kurths. Nonlinear dynamical system identification from uncertain and indirect measurements. *International Journal of Bifurcation and Chaos*, 14(06):1905–1933, 2004.
- [13] I. J. Wang and J. C. Spall. Stochastic optimisation with inequality constraints using simultaneous perturbations and penalty functions. *International Journal of Control*, 81(8):1232–1238, 2008.
- [14] Qing Wang, Sanjeev R. Kulkarni, and Sergio Verdu. Divergence estimation for multidimensional densities via  $k$ -nearest-neighbor distances. *IEEE Transactions on Information Theory*, 55(5):2392–2405, 2009.
- [15] M. T. Wyss, R. Jolivet, A. Buck, P. J. Magistretti, and B. Weber. In vivo evidence for lactate as a neuronal energy source. *J Neurosci*, 31(20):7477–85, 2011.
- [16] Feng Xu and Hanshu Ding. A new kinetic model for heterogeneous (or spatially confined) enzymatic catalysis: Contributions from the fractal and jamming (overcrowding) effects. *Applied Catalysis A: General*, 317(1):70–81, 2007.